arXiv:1402.0423v1 [cs.DS] 3 Feb 2014

A Bound on the Expected Optimality of Random Feasible Solutions to Combinatorial Optimization Problems

Evan A. Sultanik
The Johns Hopkins University APL*
evan@sultanik.com
http://www.sultanik.com/

February 4, 2014

Abstract

This paper investigates and bounds the expected solution quality of combinatorial optimization problems when feasible solutions are chosen at random. Loose general bounds are discovered, as well as families of combinatorial optimization problems for which random feasible solutions are expected to be a constant factor of optimal. One implication of this result is that, for graphical problems, if the average edge weight in a feasible solution is sufficiently small, then any randomly chosen feasible solution to the problem will be a constant factor of optimal. For example, under certain well-defined circumstances, the expected constant of approximation of a randomly chosen feasible solution to the Steiner network problem is bounded above by 3. Empirical analysis supports these bounds and actually suggest that they might be tightened.

1 Introduction

Combinatorial optimization is the process of selecting a set of points from a finite topological space that maximize or minimize a given objective function. It is often very simple to find *feasible* solutions to such problems, however, it is usually exceedingly hard to find *optimal* ones. For example, given a finite set of objects $y \in Y$ each with value $v : Y \to \mathbb{R}$ and weight $w : Y \to \mathbb{R}_{\geq 0}$, the knapsack problem asks to find a subset of the objects $S \subseteq Y$ whose combined weight does not exceed a given maximum, $w_{\max}$, and whose combined value is

*Evan is no longer affiliated with Johns Hopkins University. He is currently the chief scientist at Digital Operatives and adjunct faculty at the Drexel University Department of Computer Science.

maximized:

$$\begin{aligned} \text{maximize} \quad & \sum_{y \in Y} v(y) x_y \\ \text{subject to:} \quad & \\ & \sum_{y \in Y} w(y) x_y \leq w_{\max}, \\ & x_y \in \{0,1\}, \qquad \forall y \in Y, \end{aligned}$$

where the chosen set of objects is $S = \{y \in Y : x_y = 1\}$. Finding a feasible solution to the knapsack problem is trivial: Any set of objects $S \subseteq Y$ that does not outweigh the maximum,

$$\sum_{s \in S} v(s) \leq w_{\max},$$

will be a feasible solution. Such a set of objects can be chosen in linear time.

As a second example, consider another **NP**-Complete problem: the Steiner network problem [1]. Given an undirected, edge-weighted graph on n vertices and k edges, $G = \langle V, E \rangle$, and also given a subset of "terminal" vertices $T \subseteq V$ of cardinality $\alpha = |T|$, the Steiner network problem asks to find a minimum weight forest that spans the terminals [2]:

$$\begin{aligned} \text{minimize} \quad & \sum_{e \in E} w(e) x_e \\ \text{subject to:} \quad & \\ & \sum_{e \in \delta(S)} x_e \geq f(S), \quad \forall S \subset V : S \neq \emptyset \\ & x_e \geq 0, \qquad \forall e \in E, \end{aligned}$$

where $\delta(S)$ is the set of edges having exactly one endpoint in S and $f : 2^V \to \{0,1\}$ is a function such that $f(S) = 1$ if and only if $\emptyset \neq S \cap T \neq T$. Despite the fact that finding an optimal solution is hard, finding a *feasible* solution is not: Any spanning forest will be a feasible solution to the Steiner network problem.

Given that finding a feasible solution is often so simple relative to finding the optimal, we would like to know: In expectation, how much worse is a given feasible certificate compared to the optimal solution? This paper investigates and bounds the expected solution quality to such problems when feasible solutions are chosen at random.

We proceed by bounding the expected value of the best- and worst-case outcomes of any choice of objects. In general, let m be a lower bound on the number of objects in an optimal solution. In the case of the Steiner network problem, for example, m must be greater than or equal to $\lfloor \frac{\alpha}{2} \rfloor$ edges. Let ℓ be an upper bound on the number of objects chosen by a given algorithm that is guaranteed to produce feasible—but not necessarily optimal—solutions. In the worst case, the algorithm will have chosen the ℓ costliest objects while the

optimal solution actually contained the m cheapest objects. How likely is this outcome? If we can bound the ratio of the expected value of the sum of the ℓ costliest objects over the expected value of the sum of the m cheapest objects, this will be an upper bound on the expected constant of approximation of the algorithm, *regardless* of how the algorithm actually chooses a feasible solution, and also regardless of which combinatorial optimization problem we are solving.

The bound on the expected constant of approximation is codified in the remainder of this paper. Section 2 serves to formalize the problem and Section 3 develops the bound, the theoretical consequences of which are discussed in Section 4. The theoretical bound is supported through empirical evaluation in Section 5. Proofs of the various lemmas and theorems are provided at the end of the paper in A.

2 Framing the Problem

It is often useful to think of combinatorial optimization as the process of selecting an optimal subset from a given finite set of objects. We proceed by treating the cost of each object in a combinatorial optimization problem as a random variable. Let $\mathbf{X} = [X_1, \ldots, X_k]$ be a vector random variable consisting of identically distributed (but not necessarily independent) random variables. In a sense, each variable captures whether or not one of the k objects from the combinatorial optimization problem will be chosen to be a part of the solution. In the knapsack problem, for example, k is the total number of objects. In the Steiner network problem, k is the number of edges in the graph. The order statistics of $\mathbf{X}$ are denoted $X_{(1:k)} \leq \ldots \leq X_{(r:k)} \leq \ldots \leq X_{(k:k)}$. For $\ell \in \{1, 2, \ldots, k\}$, let Y be the distribution of the ℓ largest order statistics of $\mathbf{X}$:

$$Y \sim \sum_{r=k-\ell+1}^{k} X_{(r:k)}.$$

Similarly, for $m \in \{1, 2, \ldots, k\}$, let Y^* be the sum of the m smallest order statistics of $\mathbf{X}$:

$$Y^* \sim \sum_{r=1}^{m} X_{(r:k)}.$$

Without loss of generality, we assume in the remainder of our analysis that the objective function is being minimized. In that case, observe that Y is the probability distribution of the cost of the costliest solution of size ℓ and, similarly, Y^* is the distribution of the cost of the best possible solution of size m. The expected value of Y will be an upper bound on the expected value of the costliest solution to the problem, whereas the expected value of Y^* will be a lower bound on the expected value of the optimal solution.

We are interested in answering the question: What can be said of the relationship between Y and Y^*? If Y and Y^* are nonnegative[1] then the expected

[1] *i.e.*, their distributions are truncated such that the probability density for all negative values is zero.

value of Y/Y^* will be an upper bound on the expected constant of approximation of any algorithm that produces feasible solutions to the optimization problem. In general, for arbitrary Y and Y^*, the expected approximation factor is bounded above by

$$1 + \left| \frac{E[Y] - E[Y^*]}{E[Y^*]} \right|. \tag{1}$$

We develop an upper bound on this expression in the following section.

3 Bounds on Order Statistics

Gascuel and Caraux discovered the following bounds on the expected value of order statistics of random variables with symmetrical distributions with mean μ and variance σ^2:

Theorem 1 (Proposition 2 of [3])**.** *The bounds*

$$\left.\begin{array}{ll} \text{if } \frac{r}{k} \leq \frac{1}{2}, & \mu - \sigma\sqrt{\frac{k}{2r}} \\ \text{if } \frac{r}{k} \geq \frac{1}{2}, & \mu - \sigma\sqrt{\frac{k(k-r)}{2r^2}} \end{array}\right\} \leq E[X_{(r:k)}] \leq \left\{\begin{array}{ll} \mu + \sigma\sqrt{\frac{k}{2(k-r+1)}} & \text{if } \frac{r-1}{k} \geq \frac{1}{2}, \\ \mu + \sigma\sqrt{\frac{k(r-1)}{2(k-r+1)^2}} & \text{if } \frac{r-1}{k} \leq \frac{1}{2}, \end{array}\right.$$

are valid and may be reached for some distributions.

With the result of Theorem 1 and some intermediary bounds on the generalized harmonic series, we are led directly to a loose lower bound for $E[Y^*]$ and a loose upper bound for $E[Y]$.

To achieve these bounds, let us first constrain the magnitude of m and ℓ with respect to k:

$$\frac{m}{k} \leq \frac{1}{2} \wedge \frac{\ell - 1}{k} \leq \frac{1}{2}.$$

These constraints appear to be reasonable: The family of combinatorial optimization problems that conform to these constraints is large. In the context of the Steiner network problem, for example, these constraints equate to requiring that both the optimal and feasible solutions encompass fewer than half of all possible edges in the graph. This assumption clearly holds for almost every simple graph since

1. such graphs have $\binom{n}{2} = \frac{n}{2}(n-1)$ possible edges;
2. any feasible solution to the Steiner network problem will have at most $n-1$ edges; and
3. $\forall n \geq 4 : n - 1 \leq \frac{n}{4}(n-1)$.

The constraints on the magnitude of m and ℓ lead to the following loose bounds:

Theorem 2 (Proof of this and the remainder of the theorems in this paper are given in A)**.**

$$\frac{m}{k} \leq \frac{1}{2} \implies E[Y^*] > m\mu - \sigma\sqrt{2k}\left(\sqrt{m+1}-1\right).$$

Theorem 3.

$$\frac{\ell-1}{k} \leq \frac{1}{2} \implies E[Y] \leq \ell\mu + \frac{\sigma\sqrt{2k}}{2}\left(2\sqrt{\ell}-1\right).$$

It is possible to bound the expected approximation factor of Equation (1) by utilizing the bounds on $E[Y^*]$ and $E[Y]$:

Theorem 4. *If* $2(\ell-1) \leq k$ *and* $2m \leq k$ *then*

$$1 + \left|\frac{E[Y]-E[Y^*]}{E[Y^*]}\right| \leq \begin{cases} \dfrac{\ell\mu - 2m\mu + 3\sigma k}{\sigma k - m\mu}, & \text{if } \mu < \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m} \\ \dfrac{\ell\mu + \sigma k}{m\mu - \sigma k}, & \text{if } \mu > \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m}, \end{cases}$$

which simplifies to

$$\leq \begin{cases} 2, & \text{if } \mu < \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m} \wedge (\ell\mu \leq -\sigma k \vee m\mu > \sigma k) \\ 3, & \text{if } \mu \leq 0 \\ 4, & \text{if } \mu(\ell+2m) \leq \sigma k \\ \varepsilon \in \mathbb{R}_{\geq 2}, & \text{if } \mu \leq \frac{\sigma k(\varepsilon-3)}{\ell-2m+\varepsilon m} < \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m} \\ \varepsilon \in \mathbb{R}_{\geq k}, & \text{if } \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m} < \mu \leq \frac{\sigma(\varepsilon-k)}{\ell}. \end{cases}$$

Corollary 1. *If*

$$\ell > \exp\left(2\sqrt{m+1}-3\right) \wedge \mu > \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m} \wedge m \leq \ell \leq \frac{k}{2},$$

then the upper bound simplifies to 3*.*

Corollary 2. *If* $\sigma \gg \mu$*, the constant bound of 3 will hold asymptotically.*

4 Discussion

Theorem 4 is a rather surprising result: If we know that the size of the optimal solution is bounded below by m and the average edge weight in a random feasible solution is small, then *any* randomly chosen solution of size at most ℓ will, on average, be a constant factor of optimal. If the edges are weighted from a distribution with a nonpositive mean, then any randomly chosen feasible solution of any size is guaranteed to be, on average, no more than three times the cost of the optimal solution.

Consider the Steiner network problem as an example. If there are α terminals then we know that the optimal solution must have at least $\left\lfloor \frac{\alpha}{2} \right\rfloor$ edges (this is m). Any feasible solution to the problem is going to be an acyclic graph, which will have at most $n-1$ edges (this is ℓ). The number of random variables, k, is equal to the number of edges, which is bounded above by $\binom{n}{2}$. If μ is small[2] then, by Theorem 4, the expected constant of approximation for any randomly chosen feasible solution to the Steiner network problem is bounded above by

$$\frac{\mu(n-1) + 3\sigma\binom{n}{2} - 2\mu\left\lfloor \frac{\alpha}{2} \right\rfloor}{\sigma\binom{n}{2} - \mu\left\lfloor \frac{\alpha}{2} \right\rfloor},$$

which quickly[3] converges to 3 as $n \to \infty$.

5 Empirical Evidence

In order to demonstrate the validity and tightness of the bound, we empirically evaluate the optimality of randomly chosen feasible solutions to combinatorial optimization problems. First, let us consider a relatively simple problem that is in **P**: minimum spanning tree. We will later also examine the **NP**-Complete Steiner network problem.

Random graphs are generated using the Erdős-Rényi model $G(n, m)$ with $n \in \{4, \ldots, 30\}$ vertices and $m \in \{3, \ldots, \binom{n}{2}\}$ edges. The edges of the random graphs are weighted according to various probability distributions (listed in the forthcoming figures, below). A random feasible solution is chosen by solving the minimum spanning tree problem on the graph while *ignoring* edge weights. While this will in effect produce a "random" feasible solution that is independent of the edge weights (and therefore also the optimization problem's objective function), note that this method is not equivalent to drawing uniformly from the set of all possible feasible solutions. As such, these results may slightly bias the random solutions' optimality. The optimal solution is discovered by solving the problem taking edge weights into account. The constant of approximation is then calculated by taking the ratio of the cost of the random feasible solution over the cost of the optimal solution. As we can see in Figure 1, the randomly chosen solutions rapidly converge to an approximation constant of 2.

Next, let us consider the Steiner network problem. The random graphs are generated similarly to the minimum spanning tree case above. The number of terminals in the problem is set to $\frac{n}{2}$ and the terminals are chosen randomly from the vertices according to a uniform distribution. The random feasible solutions are chosen by solving the minimum spanning tree problem on the graph. Since the Steiner network problem is computationally intractable to solve optimally, though, we are forced to compare the randomly chosen feasible solutions against a bounded approximation to the optimal solution using a 2-Optimal algorithm [1]. Therefore, if the random feasible solutions have a cost no

[2] *i.e.*, the first case of the piecewise bound of Theorem 4 is satisfied.

[3] Convergence is superlinear, as confirmed by d'Alembert's ratio test.

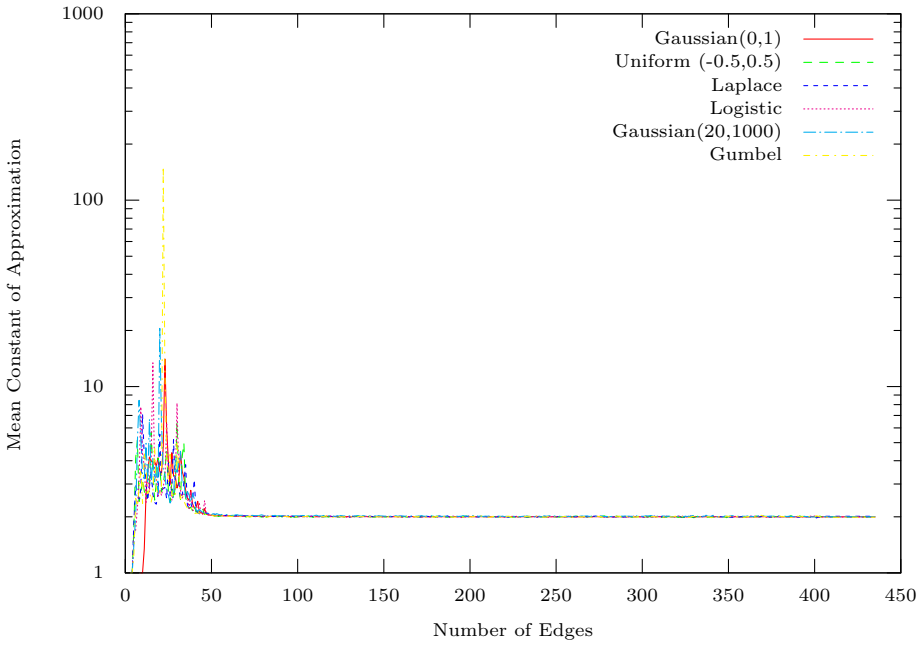

Figure 1: The average constant of approximation of randomly chosen feasible solutions to the minimum spanning tree problem converges to 2.

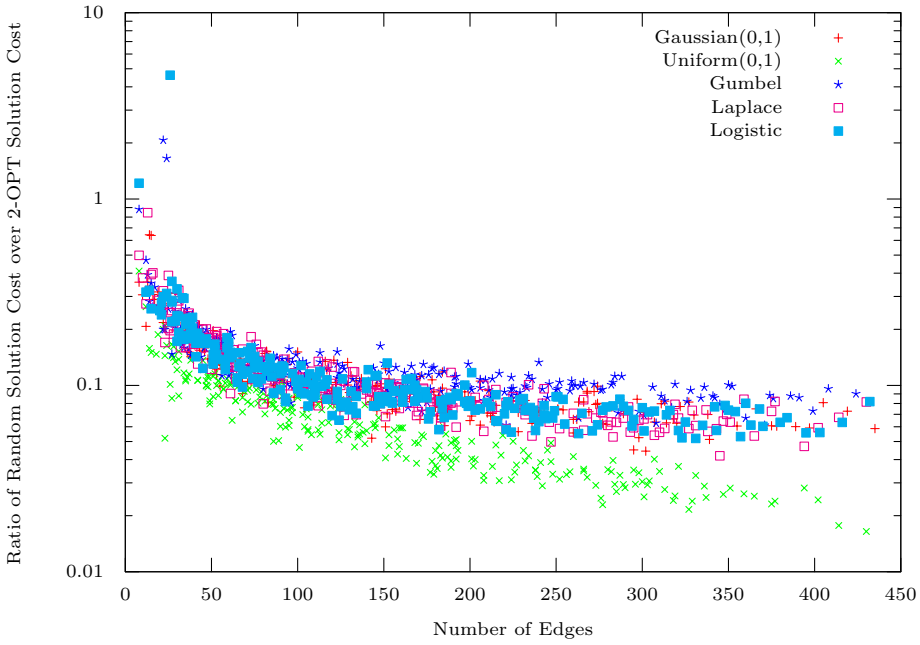

Figure 2: Random feasible solutions to the Steiner network problem compared to a 2-OPT algorithm. Values less than 1 imply that the random solution was lower cost than the approximation algorithm's solution.

worse than the approximation algorithms, we know that the cost of the feasible solutions are no more than twice the cost of the optimal solution. As we see in Figure 2, for almost all graphs the randomly chosen solutions actually perform better than the approximation algorithm, with improvement as the graphs get denser.

6 Conclusions

This paper investigated and bounded the expected solution quality of combinatorial optimization problems when feasible solutions are chosen at random. A loose general bound was discovered (Theorem 4), as well as families of combinatorial optimization problems for which random feasible solutions are expected to be a constant factor of optimal. One implication of this result is that, for graphical problems, if the average edge weight in a feasible solution is sufficiently small, then any randomly chosen feasible solution to the problem will be a constant factor of optimal. For example, under certain well-defined circumstances, the expected constant of approximation of a randomly chosen feasible solution to the Steiner network problem is bounded above by 3. Empirical analysis supports these bounds and actually suggest that they might be tightened.

Acknowledgements

This research was performed under The Johns Hopkins University Applied Physics Laboratory research grant X9RSTHNA and the Stuart S. Janney award. Thanks are owed to Kevin Schultz and Ali Shokoufandeh for their helpful comments and insight.

A Proofs and Supplementary Theorems

Let $H_{n,r}$ denote the generalized harmonic number of order n of r:

$$H_{n,r} = \sum_{i=1}^{n} \frac{1}{i^r}.$$

Lemma 1.

$$2\sqrt{n+1} - 2 < H_{n,\frac{1}{2}} \leq 2\sqrt{n} - 1.$$

Proof. Since $\frac{1}{\sqrt{x}}$ is monotonically decreasing in x, we can bound $H_{n,\frac{1}{2}}$ by its definite integral using the forward and backward rectangle rules. For the lower bound,

$$\begin{aligned} H_{n,\frac{1}{2}} &= \sum_{i=1}^{n} \frac{1}{\sqrt{i}} \\ &> \int_{1}^{n+1} \frac{1}{\sqrt{x}}\, dx \\ &= 2\sqrt{n+1} - 2. \end{aligned}$$

For the upper bound,

$$\begin{aligned} H_{n,\frac{1}{2}} &= 1 + \sum_{i=2}^{n} \frac{1}{\sqrt{i}} \\ &\leq 1 + \int_{1}^{n} \frac{1}{\sqrt{x}}\, dx \\ &= 2\sqrt{n} - 1. \end{aligned}$$

□

Proof of Theorem 2. The conditions of the implication in the theorem ensure that the bounds of Theorem 1 simplify to

$$E[X_{(r:k)}] \geq \mu - \sigma\sqrt{\frac{k}{2r}}.$$

Applying this bound to the trimmed sum of the m smallest order statistics results in

$$m\mu - \sum_{r=1}^{m} \sigma\sqrt{\frac{k}{2r}}.$$

Further simplification yields

$$m\mu - \frac{1}{2}\sigma\sqrt{2k}\sum_{r=1}^{m}\sqrt{\frac{1}{r}}. \tag{2}$$

By Lemma 1, (2) is bounded below by

$$m\mu - \sigma\sqrt{2k}\left(\sqrt{m+1}-1\right).$$

□

Proof of Theorem 3. The conditions of the implication in the theorem ensure that the bounds of Theorem 1 simplify to

$$E[X_{(r:k)}] \leq \mu + \sigma\sqrt{\frac{k}{2(k-r+1)}}.$$

Applying this bound to the trimmed sum of the ℓ largest order statistics results in

$$\ell\mu + \sum_{r=k-\ell+1}^{k}\sigma\sqrt{\frac{k}{2k-2r+2}}.$$

Further simplification yields

$$\ell\mu + \frac{1}{2}\sigma\sqrt{2k}\sum_{r=k-\ell+1}^{k}\sqrt{\frac{1}{k-r+1}}. \tag{3}$$

The summation in (3) is equal to $H_{\ell,\frac{1}{2}}$:

$$\begin{aligned}\sum_{r=k-\ell+1}^{k}\sqrt{\frac{1}{k-r+1}} &= \sqrt{\frac{1}{\ell}}+\sqrt{\frac{1}{\ell-1}}+\sqrt{\frac{1}{\ell-2}}+\ldots+\sqrt{\frac{1}{1}}\\ &= \sum_{i=1}^{\ell}\frac{1}{\sqrt{i}}\\ &= H_{\ell,\frac{1}{2}},\end{aligned}$$

which, by Lemma 1, is bounded above by $2\sqrt{\ell}-1$. □

Proof of Theorem 4. Let $\overline{E[Y]}$ be the upper bound of Theorem 3 and $\underline{E[Y^*]}$ be the lower bound of Theorem 2. Note that

$$1+\left|\frac{E[Y]-E[Y^*]}{E[Y^*]}\right| \leq 1+\left|\frac{\overline{E[Y]}-\underline{E[Y^*]}}{\underline{E[Y^*]}}\right|,$$

allowing us to proceed by bounding the expression using $\overline{E[Y]}$ and $\underline{E[Y^*]}$. Next, observe that $\overline{E[Y]} \geq \underline{E[Y^*]}$, so the absolute value in the numerator can be dropped:

$$1 + \left| \frac{E[Y] - E[Y^*]}{E[Y^*]} \right| \leq 1 + \frac{\overline{E[Y]} - \underline{E[Y^*]}}{\left| \underline{E[Y^*]} \right|}.$$

There are therefore two cases to consider, depending on the sign of the denominator. Let us first consider the case when $\underline{E[Y^*]} < 0$. Since m, k, and σ are necessarily positive, observe that

$$\underline{E[Y^*]} < 0 \iff \mu < \sigma\sqrt{2k}\frac{\sqrt{m+1}-1}{m},$$

producing the first case in the piecewise function of the theorem. Under these circumstances, evaluating Equation 1 yields

$$\begin{aligned} 1 + \left| \frac{E[Y] - E[Y^*]}{E[Y^*]} \right| &\leq 1 + \frac{\overline{E[Y]} - \underline{E[Y^*]}}{\underline{E[Y^*]}} \\ &= \frac{\ell\mu - 2m\mu + \sigma\sqrt{2k}\left(\sqrt{\ell} + 2\sqrt{m+1} - \frac{5}{2}\right)}{\sigma\sqrt{2k}\left(\sqrt{m+1} - 1\right)m\mu}. \quad (4) \end{aligned}$$

Since $m \geq 1 \implies \sigma\sqrt{2k}\left(\sqrt{m+1} - 1\right) \geq \sigma k$ and

$$\ell, m \geq 1 \implies \sigma\sqrt{2k}\left(\sqrt{\ell} + 2\sqrt{m+1} - \frac{5}{2}\right) \leq 3k,$$

the bound can be relaxed to

$$(4) \leq \frac{\ell\mu - 2m\mu + 3\sigma k}{\sigma k - m\mu}.$$

In the second and final case, $\underline{E[Y^*]} > 0$ and

$$\begin{aligned} 1 + \left| \frac{E[Y] - E[Y^*]}{E[Y^*]} \right| &\leq 1 + \frac{\overline{E[Y]} - \underline{E[Y^*]}}{\underline{E[Y^*]}} \\ &= \frac{\ell\mu + \sigma\sqrt{2k}(\sqrt{\ell} - \frac{1}{2})}{m\mu - \sigma\sqrt{2k}(\sqrt{m+1} - 1)} \leq \frac{\ell\mu + \sigma k}{m\mu - \sigma k}, \end{aligned}$$

because $2(\ell-1), 2m \leq k$ implies $\sqrt{2k}(\sqrt{\ell}+2\sqrt{m+1}-\frac{5}{2}) \leq k$ and $\sqrt{2k}(\sqrt{m+1}-1) \leq k$. □

Proof of Corollary 1. Evaluating the expression yields

$$\begin{aligned}
\frac{2\ell\mu + \sigma\sqrt{2k}\,(1+\log\ell)}{2m\mu + 2\sigma\sqrt{2k}\left(\sqrt{m+1}-1\right)} &\leq \frac{\ell}{m} \\
&\Downarrow \\
2\mu + \frac{\sigma\sqrt{2k}\,(\log\ell+1)}{\ell} &\leq 2\mu + \frac{2\sigma\sqrt{2k}\left(\sqrt{m+1}-1\right)}{m} \\
&\Downarrow \\
\frac{\log\ell+1}{\ell} &\leq \frac{2\left(\sqrt{m+1}-1\right)}{m} \\
&\Downarrow \\
\ell &\leq \exp(2\sqrt{m+1}-3),
\end{aligned}$$

which is satisfied by the conditions of the corollary. □

Proof of Corollary 2. From Theorem 4,

$$\sigma \gg \mu \implies 1 + \left|\frac{E[Y]-E[Y^*]}{E[Y^*]}\right| \leq \frac{\ell\mu - 2m\mu + 3\sigma k}{\sigma k - m\mu} = \Theta(3).$$

□

References

[1] M. X. Goemans, A. V. Goldberg, S. Plotkin, D. B. Shmoys, E. Tardos, D. P. Williamson, Improved approximation algorithms for network design problems, in: Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms, 1994, pp. 223–232.

[2] M. Aggarwal, N. Garg, A Scaling Technique for Better Network Design, in: Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1994, pp. 233–240.

[3] O. Gascuel, G. Caraux, Bounds on expectations of order statistics via extremal dependences, Statistics & Probability Letters 15 (2) (1992) 143–148.